REPRESENTATIONAL DISTANCE LEARNING FOR DEEP NEURAL NETWORKS

Patrick McClure & Nikolaus Kriegeskorte *
MRC Cognition and Brain Science Unit
Cambridge, UK
{Patrick.McClure,Nikolaus.Kriegeskorte}@mrc-cbu.cam.ac.uk

arXiv:1511.03979v5 [cs.NE] 4 Dec 2015

ABSTRACT

We propose representational distance learning (RDL), a technique that allows transferring knowledge from an arbitrary model with task related information to a deep neural network (DNN). This method seeks to maximize the similarity between the representational distance matrices (RDMs) of a model with desired knowledge, the teacher, and a DNN currently being trained, the student. The knowledge contained in the information transformations performed by the teacher are transferred to a student using auxiliary error functions. This allows a DNN to simultaneously learn from a teacher model and learn to perform some task within the framework of backpropagation. We test the use of RDL for knowledge distillation, also known as model compression, from a large teacher DNN to a small student DNN using the MNIST and CIFAR-10 datasets. Also, we test the use of RDL for knowledge transfer between tasks using the CIFAR-10 and CIFAR-100 datasets. For each test, RDL significantly improves performance when compared to traditional backpropagation alone and performs similarly to, or better than, recently proposed methods for model compression and knowledge transfer.

1 INTRODUCTION

Deep neural networks (DNNs) have recently been highly successful for machine perception, particularly in the areas of computer vision using convolutional neural networks (CNNs) (Krizhevsky et al., 2012) and speech recognition using recurrent neural networks (RNNs) (Deng et al., 2013). The success of these methods depends on their ability to learn good, hierarchical representations for these tasks (Bengio, 2012).

Traditionally, propagation of the output error signal from the last layer of a network back to the first layer, known as backpropagation, has been the main method for learning these hierarchical representations (Rumelhart et al., 1988). Recently, it has been investigated how the error signal reaching an internal layer through backpropagation can be complemented by auxiliary error functions that more directly constrain the representation. A variety of criteria have successfully been used as sources of auxiliary information for several DNNs.

Weston et al. (2012) proposed semi-supervised embedding to augment the error from the output layer. A reference embedding of the inputs was used to guide representational learning. The embedding constraint was implemented in different ways: inside the network as a layer, as part of the output layer, or as an auxiliary error function that directly affected a particular hidden layer. Weston et al. discussed a variety of embedding methods that could be used, including multidimensional scaling (MDS) (Kruskal, 1964) and Laplacian Eigenmaps (Belkin & Niyogi, 2003). Ultimately, they used a LapSVM (Belkin et al., 2006) auxiliary error, which is a linear combination of hinge loss (L2SVM) (Rosasco et al., 2004) and Laplacian Eigenmaps with L2 parameter regularization. The addition of these semi-supervised error functions led to increased accuracy compared to DNNs trained using output layer backpropagation alone.

Lee et al. (2014) also showed that auxiliary error functions improve DNN representational learning. Instead of using semi-supervised methods, they performed classification with a softmax or L2SVM

*Corresponding author

readout at a given intermediate hidden layer. The error of the intermediate-level readout was then backpropagated to earlier layers. The gradients from these classifiers were linearly combined with the gradients from the output layer classifier. This technique resulted in state-of-the-art accuracies for several datasets.

A challenge on training very deep networks is the problem of vanishing gradients. Layers far from the output may receive only a weak learning signal via conventional backpropagation. Auxiliary error functions were successfully applied by Szegedy et al. (2014) to inject a complementary learning signal at internal layers by constraining representations to better discriminate between classes. This was implemented in a very large CNN which won the ILSVRC14 classification competition (Russakovsky et al., 2014). In this DNN, two auxiliary networks were used to directly backpropagate from two intermediate layers back through the main network. Each auxiliary network contained one convolutional layer, two fully connected layers, and a softmax classifier. Similar to the method used in Lee et al. (2014), the parameters for the layers in the main network directly connected to auxiliary networks were updated using a linear combination of the backpropagated gradients from later layers and the auxiliary network.

Wang et al. (2015) investigated the effectiveness of auxiliary error functions in very large CNNs and their optimal placement. They selected where to place these auxiliary functions by measuring the average magnitude of the conventional backpropagation error signal at each layer. Auxiliary networks, similar to those used in Szegedy et al. (2014), were placed after layers with vanishing gradients. These networks consisted of a convolutional layer followed by three fully connected layers and a softmax classifier. As in Lee et al. (2014) and Szegedy et al. (2014), the auxiliary gradients were linearly combined to update the model parameters. Adding these supervised auxiliary error functions led to an improved accuracy for two very large datasets, ILSVRC12 (Russakovsky et al., 2014) and MIT Places (Zhou et al., 2014).

In the studies just cited, auxiliary error functions served to pull representations toward the desired output. Auxiliary error functions can also be used learn representations resembling those of other models. This is useful for model compression (also known as knowledge distillation) and transfer learning (Bengio, 2012). The goal in either case is to use the representational knowledge learned by a teacher neural network to improve the performance of a student network (Bucilua et al., 2006; Ba & Caruana, 2014; Hinton et al., 2015). For model compression, the teacher is a larger or more complex network with higher performance than the student. For knowledge transfer, the representations learned by the teacher network are used to improve the training of a student network on a different tasks or using different data. Several methods have been proposed to perform these tasks.

One technique for model compression is to have the student learn the output representation of the teacher for a given training input. For classification, the neurons before the softmax layer can be constrained to have the same values as the teacher using mean squared error (MSE) as done in Bucilua et al. (2006); Ba & Caruana (2014). Alternatively, the output of the softmax layer can be constrained to represent the same, or similar, output distribution as the teacher. This can be done by minimizing the cross-entropy between the output distributions of the teacher and student networks for the training inputs (Hinton et al., 2015). Hinton et al. used the softmax output of a large teacher network trained using dropout (Srivastava et al., 2014) to train a student using soft targets. Dropout is a regularization technique that approximates model averaging by turning off neurons in a layer with a probability of p during training and then multiplying the output of the layer by p during testing. To form soft targets, the inputs to the softmax layer of a large dropout network were divided by a temperature parameter before softmax was applied. These outputs where then used as soft targets in addition to class labels to train a student network to approximate the teacher network.

Knowledge from different networks can also be transferred at internal layers. Romero et al. (2014) proposed a method for transferring the knowledge of a wide and shallow teacher to a thin and deep student. This technique constrained an intermediate layer of the student network to have representations that could linearly predict those of the teacher network. This method was shown to improve the students classification accuracy.

One prominent technique for performing transfer learning is to initialize the weights of the student network to those of the teacher. The network is then trained on a different task or using different data. This can lead to improved network performance (Yosinski et al., 2014). However, this requires that the teacher and student have the same, or very similar, architectures, which may not be desirable.

In this paper, we introduce an auxiliary error function that enables a student network to learn from the internal representational spaces of a teacher that has a different architecture. The method constrains the student's representational distances in internal layers to approximate those of the teacher. The student can thus learn the computational transformations discovered by the teacher. This approach improves representational learning for classification tasks. However, unlike output-based auxiliary functions and learning from soft-targets, this technique can be used for both model compression and transfer learning.

2 METHODS

Our method, representational distance learning (RDL), allows DNNs to learn from the representations of other models to improve performance. As in Lee et al. (2014); Szegedy et al. (2014); Wang et al. (2015), we utilize auxiliary error functions to train internal layers directly in conjunction with the error from the output layer found via backpropagation. We propose an error function that maximizes the similarity between the representational spaces of a student DNN and that of a teacher model.

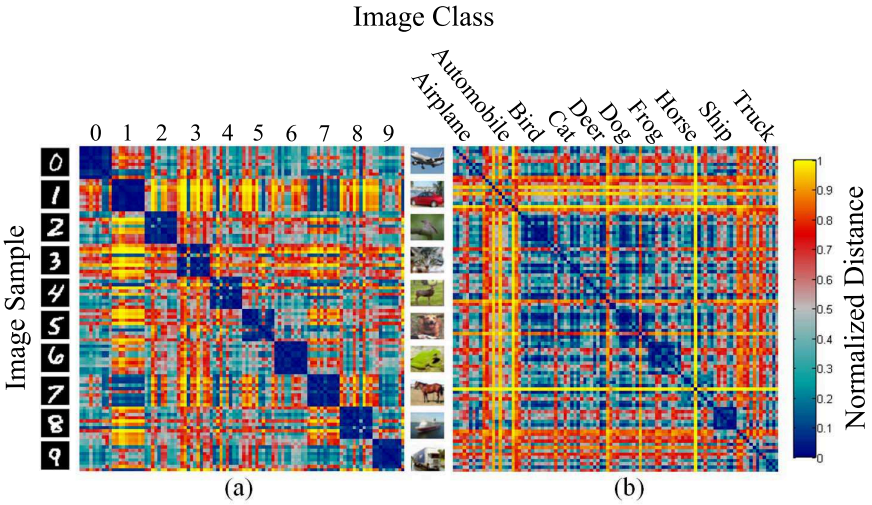

Figure 1: Example representational distance matrices (RDMs) of the output layer of convolutional neural networks (CNNs) for ten random images of each class from (a) MNIST and (b) CIFAR-10 made using the RSA toolbox (Nili et al., 2014).

2.1 REPRESENTATIONAL DISTANCE MATRICES

In order to compare the representational spaces of models, a method must be used to describe them. As discussed in Weston et al. (2012), a representational space can be characterized by the pairwise distances between representations. This idea has been used in several methods such as MDS, which seeks to reduce the dimensionality of data while minimizing the error between the pairwise distance matrix of the original data and the reduced dimensionality data (Kruskal, 1964).

Kriegeskorte et al. (2008) proposed using the matrix of pairwise dissimilarities between representations of different inputs, which they called representational distance, or dissimilarity, matrices (RDMs), to compare computational models and neurological data. More recently, Khaligh-Razavi & Kriegeskorte (2014) used this technique to analyze several computer vision models, including the CNN from Krizhevsky et al. (2012), and neurological data. Any distance function could be used to compute the pairwise dissimilarities, for instance the Euclidean or correlation distances. An RDM for a DNN can be defined by:

$$RDM(X; f_m)_{i,j} = d(f_m(x_i; W_m), f_m(x_j; W_m)) \tag{1}$$

where X is a set of n inputs (e.g. a mini-batch or a subset of a mini-batch), f_m is the neuron activations at layer m, x_i and x_j are single inputs, W_m is the weights of the neural network up to layer m, and some distance, or dissimilarity, measure d.

In addition to characterizing the information present in a particular layer of a DNN, RDMs can be used to visualize the representational space at different layers within a DNN (Figure 1). Currently, understanding and visualizing the information captured by internal layers in a DNN is challenging. Zeiler & Fergus (2014) recently proposed a method for visualizing the input features which active internal neurons at varying layers using deconvolutional neural networks. Yosinski et al. (2015) also proposed methods for visualizing the activations of a DNNs for a given input. However, these methods do not show the categorical information of each representational layer. Visualizing the similarity of labelled inputs at layers of interest, via an RDM, allow clusters inherent to the learned representation spaces to be shown. (See Section A in the Supplementary Material for more details.)

2.2 Representational Distance Learning

RDL uses auxiliary error functions to maximize the similarity of the RDMs of a student model to a teacher model at several layers. This is motivated by the idea that RDMs, or distance matrices in general, can characterize the representational space of a model. DNNs seek to learn a set of hierarchical representations. For classification, this culminates in finding a representational space where different classes are separable. RDL allows a DNN to learn from the representations of a different, potentially better, model by maximizing the similarity between the RDMs of the DNN being trained and the target model at several layers. Unlike in Bucilua et al. (2006); Ba & Caruana (2014); Hinton et al. (2015). RDL not only directly trains the output representation, but also the representations of hidden layers. As discussed in Bengio (2012), however, large datasets can prohibit the use of pairwise techniques, since the number of comparisons grows quadratically with dataset size. To address this, our technique only uses a random subset of all pairwise distances for each parameter update. This allows the speed of our method to be constrained by the subset size and not the overall number of training examples, which is usually several orders of magnitude larger.

In order to maximize the similarity between the RDM of a DNN layer being trained and a target RDM, we propose minimizing the mean squared error between the two RDMs. This corresponds to making all possible pairwise distances as similar as possible:

$$E_{aux}(X; f_m; T) = \frac{1}{2(n^2 - n)} \sum_{i=1}^{n} \sum_{j \neq i}^{n} (RDM(X; f_m)_{i,j} - T_{i,j})^2 \tag{2}$$

where X is a set of n inputs (e.g. a mini-batch or a subset of a mini-batch), f_m is the neuron activations at layer m, and $T_{i,j}$ is the target distance from input x_i and input x_j, which in the context of model compression is the teachers RDM. The function d used to calculate the RDMs (Eq. 1) could be any dissimilarity or distance function, but we chose to use the mean squared error (MSE). This results in the average auxiliary error with respect to neuron k of f_m , $f_{m,k}$, for input x_i and the weights of the neural network up to layer m, W_m, being defined as:

$$\frac{\partial E_{aux}(x_i; X; f_m; T)}{\partial f_{m,k}} = \frac{4}{n^2 - n} \sum_{j \neq i}^{n} (RDM(X; f_m)_{i,j} - T_{i,j})(f_{m,k}(x_i; W_m) - f_{m,k}(x_j; W_m)) \tag{3}$$

However, calculating the error for every pairwise distance can be computational expensive, so we estimate the error using a random subset, P, of the pairwise distances for each update of a network's parameters. This leads to the auxiliary error gradient being approximated by:

$$\frac{\partial E_{aux}(x_i; X; f_m; T)}{\partial f_{m,k}} \approx z_i \sum_{(i,j) \in P_{x_i}} (RDM(X; f_m)_{i,j} - T_{i,j})(f_{m,k}(x_i; W_m) - f_{m,k}(x_j; W_m)) \tag{4}$$

Table 1: % classification error for the MNIST, CIFAR-10, and CIFAR-100 test sets using the baseline, baseline with dropout (Dropout), baseline with soft-targets (Soft), baseline with RDL, baseline with fine-tuning and teacher networks.

Dataset	Baseline	Dropout	Soft	Fine-tuning	RDL	Teacher
MNIST	0.87	0.74	0.58	-	0.69	0.56
CIFAR-10	24.26	20.58	23.64	-	23.25	19.36
CIFAR-100	48.08	-	-	47.48	46.46	-

where $z_i = 4/(|X_P||P_{x_i}|)$, X_P is the set of all images contained in P, P_{x_i} is the set of all pairs, (i, j), in P that include input x_i and another input, x_j. If an image is not sampled, its auxiliary error is zero.

The total error of $f_{m,k}$ for input x_i is calculated by taking a linear combination of the auxiliary error at layer m and the error from backpropagation of the output error function and any later auxiliary functions. These terms are combined using weighting hyper parameter α, similar to the method discussed in Lee et al. (2014), Szegedy et al. (2014), and Wang et al. (2015). In RDL, α is the weight of the RDL error in the overall error function. Subsequently, the error gradient at a layer with an auxiliary error function is defined as:

$$\frac{\partial E_{total}(x_i; y_i; X; f_m; T)}{\partial f_{m,k}} = \frac{\partial E_{backprop}(x_i; y_i; f_m)}{\partial f_{m,k}} + \alpha \frac{\partial E_{aux}(x_i; X; f_m; T)}{\partial f_{m,k}} \tag{5}$$

This error is then used to calculate the error of earlier layers in the DNN using backpropagation. As discussed in Lee et al. (2014) and Wang et al. (2015), the value of α was decayed as training progressed. Throughout training, α was updated following $\alpha_{t+1} = \alpha_0 * (1 - t/t_{max})$ where t is the epoch number and t_{max} is the total number of epochs. By using this decay rule, the auxiliary error function initially helps drive the parameters to good values while allowing the DNN to converge predominantly using the output error by the end of training.

3 RESULTS

To evaluate the effectiveness of RDL, we performed tests using three different datasets, MNIST, CIFAR-10, and CIFAR-100. For MNIST and CIFAR-10, we trained a large teacher CNN. For both datasets, this network was then compared to a small baseline network with and without dropout trained with standard output error-based backpropagation, a small network trained using soft-targets from the teacher, and a small network trained with RDL using the RDMs of the teacher. For CIFAR-100, we tested using RDL for knowledge transfer between tasks. We attempted to transfer knowledge from a network trained to perform CIFAR-10 classification to a network being trained to perform CIFAR-100 classification using both fine tuning (Yosinski et al., 2014) and RDL. When RDL was used, auxiliary error functions were placed after each max pooling layer and before the softmax layer, similar to their placement in Lee et al. (2014), and 200 image pairs were sampled for each mini-batch update. For all experiments, stochastic gradient descent (SGD) with a momentum of 0.9, a mini-batch size of 100, and ReLU activation functions were used. For all networks with dropout, $p = 0.5$. All soft-targets were created using a temperature of 20, as done in Hinton et al. (2015). The networks trained for each task were compared using the exact McNemar test (Edwards, 1948).

3.1 MNIST

MNIST is a dataset of 28x28 images of handwritten digits from ten classes, 0 through 9 (LeCun et al., 1998). The dataset was split into 40,000 training images, 10,000 validation images, and 10,000 test images. No pre-processing or data augmentation was applied. The teacher network consisted of a 32 channel convolutional layer with 5x5 filters and max pooling, a 64 channel convolutional layer with 5x5 filters and max pooling, and, finally, a fully connected layer with 500 units and with dropout. All of the student networks had an architecture with half of the neurons as the teacher. This resulting in them having a 16 channel convolutional layer with 5x5 filters and max

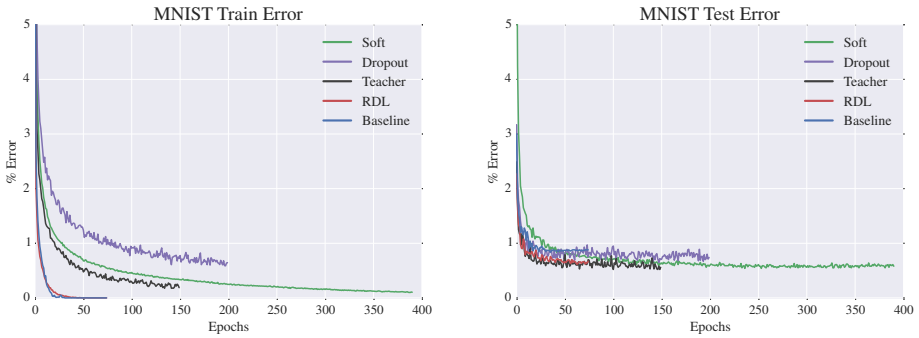

Figure 2: The change in the train and test errors through time as the baseline, baseline with dropout (Dropout), baseline with soft-targets (Soft), baseline with RDL, and teacher networks are trained on MNIST.

pooling, a 32 channel convolutional layer with 5x5 filters and max pooling, and, finally, a fully connected layer with 250 units. For RDL, we found that α parameters with a magnitude of the order of 1e-5 worked well.

The RDL network was significantly more accuate than the baseline network and more accurate than the dropout network, but not significantly. While having a lower accuracy than the soft-target and teacher networks, RDL's performance was not significantly different from these networks.

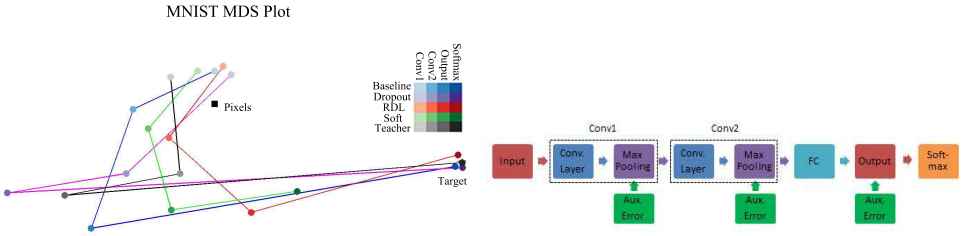

(a) (b)

Figure 3: (a) Multi-dimensional scaling (MDS) visualization of the distances between the representational distance matrices (RDMs) for selected layers of the MNIST networks with different colored lines correspond to different networks and points with darker hues correspond to deeper layers (Figure A.1). (b) A visualization of the architecture used for RDL on MNIST.

3.2 CIFAR-10

CIFAR-10 is a dataset of 32x32 color images each containing one of ten objects. The dataset was split into 40,000 training images, 10,000 validation images, and 10,000 test images. These images were pre-processed using global contrast normalization and ZCA whitening, but no data augmentation was performed. The teacher network consisted of three 128 channel convolutional layer with 5x5 filters and 3x3 max pooling then a fully connected layers each with 1000 unit and with dropout. All of the student networks had an architecture of three 64 channel convolutional layer with 5x5 filters and 3x3 max pooling then a fully connected layers each with 500 units. For RDL, we found that α parameters with a magnitude of the order of 1e-4 worked well.

Both the accuracy (Table 1) and the McNemar results (Table B.2) showed that RDL is effective for DNN compression using CIFAR-10. RDL was significantly more accurate than the baseline. RDL was also more accurate than the soft-target network, but not significantly. Additionally, RDL was less accurate than the dropout and teacher networks, but not significantly.

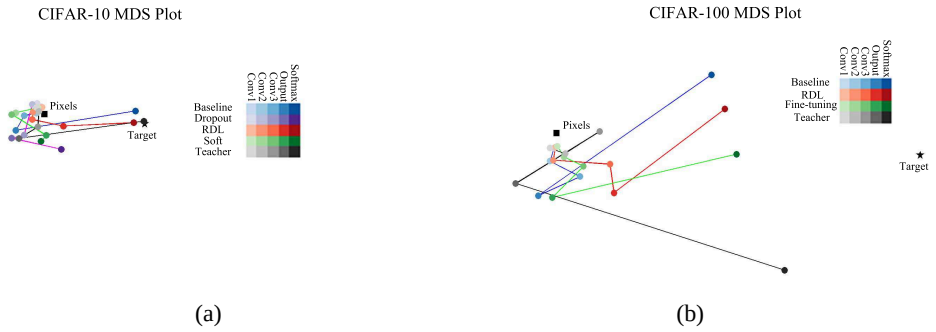

(a) (b)

Figure 4: Multi-dimensional scaling (MDS) visualization of the distances between the representational distance matrices (RDMs) for selected layers of (a) the CIFAR-10 networks (Figure A.2) and (b) the CIFAR-100 networks (Figure A.3), respectively. The different colored lines correspond to different networks and points with darker hues correspond to deeper layers.

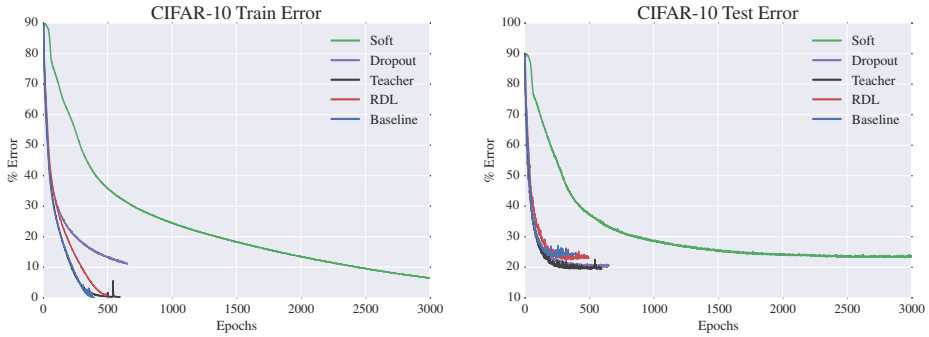

Figure 5: The change in the train and test errors through time as the baseline, baseline with dropout (Dropout), baseline with soft-targets (Soft), baseline with RDL, and teacher networks are trained on CIFAR-10.

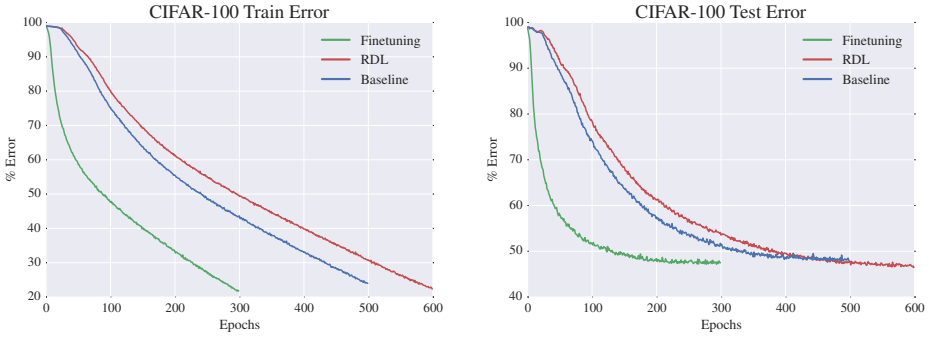

Figure 6: The change in the train and test errors through time as the baseline, baseline with finetuning, and baseline with RDL are trained on CIFAR-100.

3.3 CIFAR-100

CIFAR-100 is a dataset of 32x32 color images each containing one of 100 objects. The dataset was split into 40,000 training images, 10,000 validation images, and 10,000 test images. These images were pre-processed using global contrast normalization and ZCA whitening, but no data augmentation was performed. The baseline has the same architecture as the CIFAR-10 teacher network except with a new 100-class output layer. All layers had randomly initialized weights. The finetuning network was initialized as the CIFAR-10 teacher network with only the fully-connected layer and the new 100 unit output layer randomly initialized. The RDL network had the same

architecture as the baseline CIFAR-100 network with randomly initialized weights and the addition of auxiliary error functions that used the RDMs from the CIFAR-10 teacher. For RDL, we found that α parameters with a magnitude of the order of 1e-3 worked well.

In this experiment, RDL had the largest accuracy (Table 1). Also, the results of the McNemar exact test (Table B.3) showed that the RDL network was significantly more accurate than the baseline network and the fine-tuning network. However, the fine-tuning network was not significantly different from the baseline network. This indicates that RDL learned better representations than both the baseline and the fine-tuning networks.

4 CONCLUSIONS

In this paper, we proposed RDL, a technique for improving representational learning for DNNs. The representational space of a DNN is pulled towards that of a teacher model during training using SGD. This was performed by minimizing the difference between the pairwise distances between representations of two models at selected layers using auxiliary error functions. RDL was shown to improve the classification accuracy of baseline networks for model compression and knowledge transfer, as well as compete with, and in some experiments exceed, the performance of other recently proposed methods. Additionally, RDL converges much faster than training using soft-targets. These results show that RDL is a promising technique for representational learning. In future work, we will investigate the placement of auxiliary error functions for RDL, using RDL with larger DNNs, and RDL using non-artificial neural network models as teachers, such as brain activity data.

ACKNOWLEDGEMENTS

The authors thank Mate Lengyel, Tibor Auer, and Hunar Abdulrahman for helpful discussions and comments on a draft of the manuscript. This research was funded by the Cambridge Commonwealth, European & International Trust, the UK Medical Research Council (Program MC-A060-5PR20), and a European Research Council Starting Grant (ERC-2010-StG 261352).

REFERENCES

Ba, Jimmy and Caruana, Rich. Do deep nets really need to be deep? In *Advances in Neural Information Processing Systems*, pp. 2654–2662, 2014.

Belkin, Mikhail and Niyogi, Partha. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003.

Belkin, Mikhail, Niyogi, Partha, and Sindhwani, Vikas. Manifold regularization: A geometric framework for learning from labeled and unlabeled examples. *The Journal of Machine Learning Research*, 7:2399–2434, 2006.

Bengio, Yoshua. Deep learning of representations for unsupervised and transfer learning. *Unsupervised and Transfer Learning Challenges in Machine Learning*, 7:19, 2012.

Bucilua, Cristian, Caruana, Rich, and Niculescu-Mizil, Alexandru. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 535–541. ACM, 2006.

Deng, Li, Hinton, Geoffrey, and Kingsbury, Brian. New types of deep neural network learning for speech recognition and related applications: An overview. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pp. 8599–8603. IEEE, 2013.

Edwards, Allen L. Note on the correction for continuity in testing the significance of the difference between correlated proportions. *Psychometrika*, 13(3):185–187, 1948.

Goodfellow, Ian J, Warde-Farley, David, Mirza, Mehdi, Courville, Aaron, and Bengio, Yoshua. Maxout networks. *arXiv preprint arXiv:1302.4389*, 2013.

Hinton, Geoffrey, Vinyals, Oriol, and Dean, Jeff. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.

Khaligh-Razavi, Seyed-Mahdi and Kriegeskorte, Nikolaus. Deep supervised, but not unsupervised, models may explain it cortical representation. *PLoS Comput Biol*, 10(11):e1003915, 11 2014. doi: 10.1371/journal.pcbi.1003915. URL http://dx.doi.org/10.1371%2Fjournal.pcbi.1003915.

Kriegeskorte, Nikolaus, Mur, Marieke, and Bandettini, Peter. Representational similarity analysis–connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2, 2008.

Krizhevsky, Alex, Sutskever, Ilya, and Hinton, Geoffrey E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.

Kruskal, Joseph B. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika*, 29(1):1–27, 1964.

LeCun, Yann, Bottou, Léon, Bengio, Yoshua, and Haffner, Patrick. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Lee, Chen-Yu, Xie, Saining, Gallagher, Patrick, Zhang, Zhengyou, and Tu, Zhuowen. Deeply-supervised nets. *arXiv preprint arXiv:1409.5185*, 2014.

Nili, Hamed, Wingfield, Cai, Walther, Alexander, Su, Li, Marslen-Wilson, William, and Kriegeskorte, Nikolaus. A toolbox for representational similarity analysis. *PLoS Comput. Biol*, 10(4): e1003553, 2014.

Romero, Adriana, Ballas, Nicolas, Kahou, Samira Ebrahimi, Chassang, Antoine, Gatta, Carlo, and Bengio, Yoshua. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014.

Rosasco, Lorenzo, Vito, Ernesto De, Caponnetto, Andrea, Piana, Michele, and Verri, Alessandro. Are loss functions all the same? *Neural Computation*, 16(5):1063–1076, 2004.

Rumelhart, David E, Hinton, Geoffrey E, and Williams, Ronald J. Learning representations by back-propagating errors. *Cognitive modeling*, 5:3, 1988.

Russakovsky, Olga, Deng, Jia, Su, Hao, Krause, Jonathan, Satheesh, Sanjeev, Ma, Sean, Huang, Zhiheng, Karpathy, Andrej, Khosla, Aditya, Bernstein, Michael, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, pp. 1–42, 2014.

Srivastava, Nitish, Hinton, Geoffrey, Krizhevsky, Alex, Sutskever, Ilya, and Salakhutdinov, Ruslan. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.

Szegedy, Christian, Liu, Wei, Jia, Yangqing, Sermanet, Pierre, Reed, Scott, Anguelov, Dragomir, Erhan, Dumitru, Vanhoucke, Vincent, and Rabinovich, Andrew. Going deeper with convolutions. *arXiv preprint arXiv:1409.4842*, 2014.

Wang, Liwei, Lee, Chen-Yu, Tu, Zhuowen, and Lazebnik, Svetlana. Training deeper convolutional networks with deep supervision. *arXiv preprint arXiv:1505.02496*, 2015.

Weston, Jason, Ratle, Frédéric, Mobahi, Hossein, and Collobert, Ronan. Deep learning via semi-supervised embedding. In *Neural Networks: Tricks of the Trade*, pp. 639–655. Springer, 2012.

Yosinski, Jason, Clune, Jeff, Bengio, Yoshua, and Lipson, Hod. How transferable are features in deep neural networks? In *Advances in Neural Information Processing Systems*, pp. 3320–3328, 2014.

Yosinski, Jason, Clune, Jeff, Nguyen, Anh, Fuchs, Thomas, and Lipson, Hod. Understanding neural networks through deep visualization. *arXiv preprint arXiv:1506.06579*, 2015.

Zeiler, Matthew D and Fergus, Rob. Visualizing and understanding convolutional networks. In *Computer Vision–ECCV 2014*, pp. 818–833. Springer, 2014.

Zhou, Bolei, Lapedriza, Agata, Xiao, Jianxiong, Torralba, Antonio, and Oliva, Aude. Learning deep features for scene recognition using places database. In *Advances in Neural Information Processing Systems*, pp. 487–495, 2014.

SUPPLEMENTARY MATERIAL

A CNN RDMS

Understanding the information transoformations performed in a deep neural network (DNN), such as a convolutional neural network (CNN), using visualization can be difficult. Using representational distance matricies (RDMs) to view which inputs are represented similarly at different stages is one method for understanding the features learned by a DNN. This is often done at the output layer using confusion matrices, which show the classes that a network is having difficulty discriminating between. However, confusion matrices are computed using a classifier, which is usually only present at or after the output layer. In this section, RDMs will be used to analyse the representations learned for several networks and for several tasks.

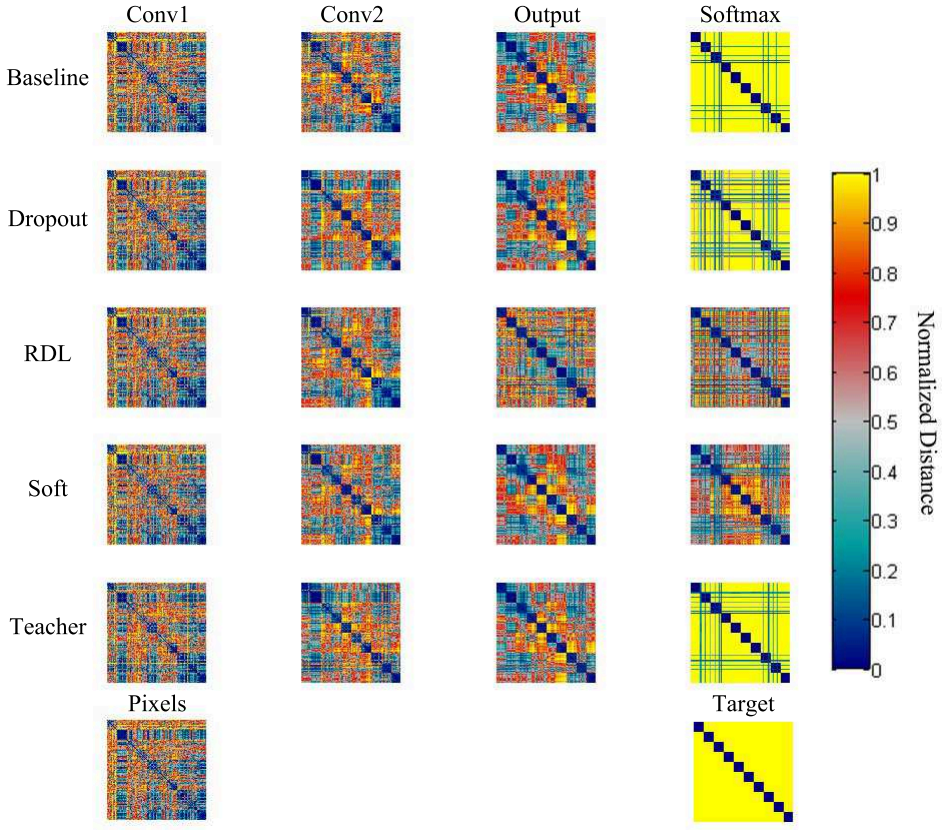

Figure 1: Representational distance matrices (RDMs) of the first and second convolutional layers as well as the output and softmax layers for the baseline, baseline with dropout (Dropout), baseline with soft-targets (Soft), baseline with RDL, and teacher networks using ten random images of each class from MNIST. Also, the RDMs of the raw pixel data and the target labels are shown.

For the MNIST trained DNNs, several pieces of information about the learned representations can be seen (Fig. A.1). For instance, images of zeros and ones are each clustered even when using only pixel intensities as features. Also, sevens and nines are represented similarly in pixel space. In general, the separability of classes does not drastically improve after conv1. However, after conv2 each class becomes more distinct, but twos are still not clustered. This changes at the output layer where the majority of inputs for each class are grouped together. As expected, the softmax layer drastically sharpens the distinction between classes for the networks trained only using output error (i.e. the baseline, dropout, and teacher networks). Even after this, the representations of a few ones that are similar to sevens and twos that are similar to eights. This is to be expected given the similarity between the classes in input space. The softmax, however, does not sharpen the class divisions as sharply for the RDL and soft-target trained networks. This seems to be due to the fact

that RDL is minimizing multiple error functions and that the soft-targets explicitly inhibit having all of the probability mass of the softmax outputs at one particular class.

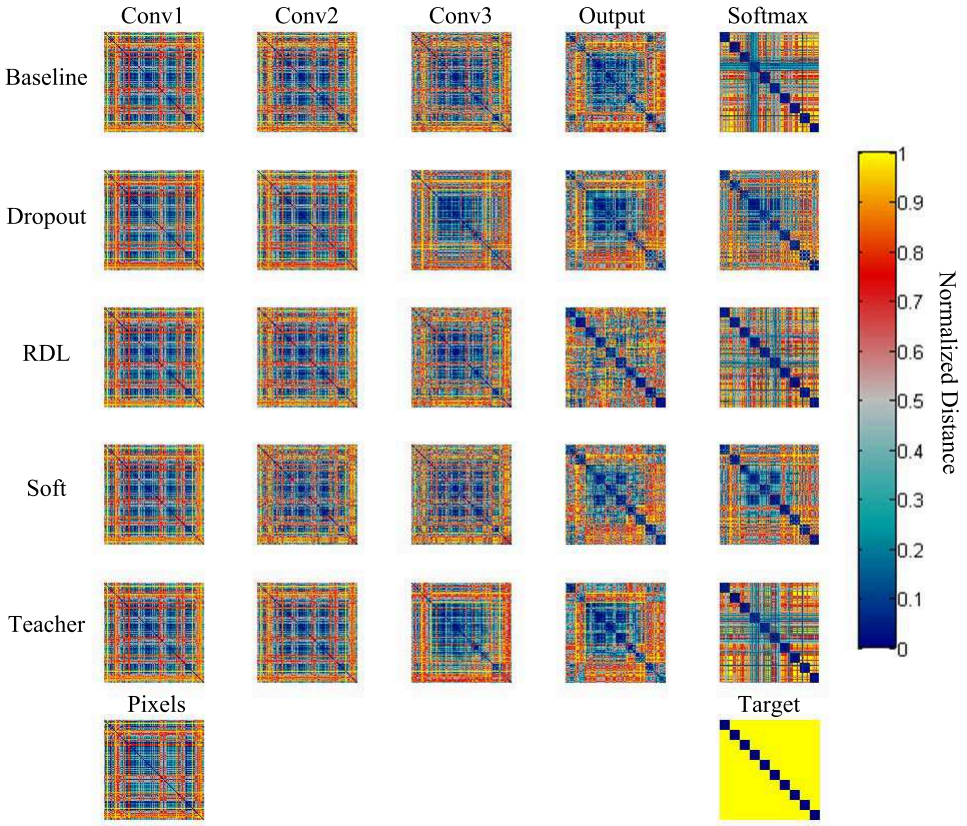

Figure 2: Representational distance matrices (RDMs) of the first, second, and third convolutional layers as well as the output and softmax layers for the baseline, baseline with dropout (Dropout), baseline with soft-targets (Soft), baseline with RDL, and teacher networks using ten random images of each class from CIFAR-10. Also, the RDMs of the raw pixel data and the target labels are shown.

The visualizations of the RDMs for the CIFAR-10 trained DNNs is also informative (Fig. A.2). The represntational space before and after conv1 and cov2 are very similar. In this space, the automobiles, horses, and trucks are very different from every other class and from each other. Also, many of the animal inputs are represented similarly. After conv3, the birds, cats, deer, dogs, and frogs are clustered together. Automobiles, horses, and trucks remain somewhat distinct. At the output layer, several classes become clustered. In the baseline, dropout, soft, and teacher networks, the representations of cats and dogs become similar to each other. Also, ships and trucks each become clustered. For the RDL network, each class becomes noticeably distinct. A similar, but weaker, effect is also present at the soft-target network output layer. As expected, the class distinctions in each network become much more apparent after the softmax layer. In the baseline and teacher networks, however, cats become somewhat similar to all classes. Cats and dogs also close in representational space for after the softmax of the dropout and soft-target networks. These classes remain pretty distinct in the RDL network, but dogs become slightly similar to many other classes.

For the RDMs of the CIFAR-100 networks (Fig. A.3), class specific is harder to observe visually. This is most likely due to the high error of the networks on the training task as well as the increased number of classes and the reduced number of samples per class. However, other characteristics of the networks' learned representations are shown. Even though the teacher network was trained on a different task, the different classes are clearly clustered. This verifies that the teacher network contains useful information for CIFAR-100 classification. Also, the RDMs show differences in the representations learned by the fine-tuning and RDL networks. The fine-tuning convolutional layers

have much better class-based clustering, similar to the teacher network, when compared to RDL. Classes are similarly distinct at the output layer, but they are more differentiated in the softmax representations for RDL. The quality of the early features learned using fine-tuning supports the hypothesis that early layers contain more general features and are to some degree task indifferent (Yosinski et al., 2014). The difference in class distinction at the softmax layer is likely because fine-tuning does not actively constrain the output layer using information from the teacher, unlike RDL.

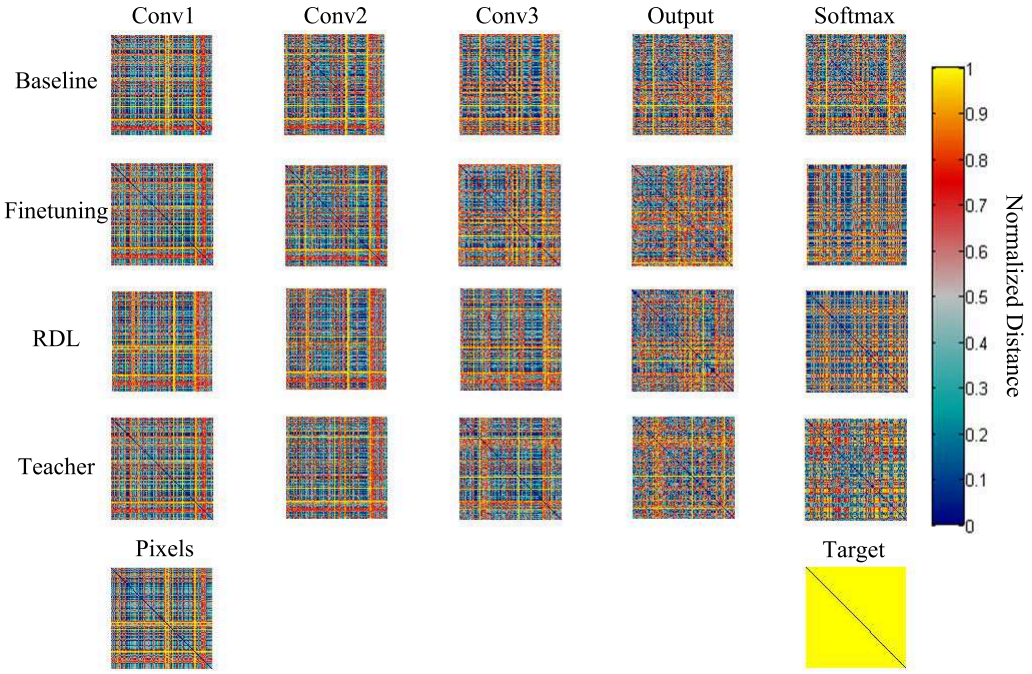

Figure 3: Representational distance matrices (RDMs) of the first, second, and third convolutional layers as well as the output and softmax layers for the baseline, baseline with fine-tuning, baseline with RDL, and teacher networks using three random images of each class from CIFAR-100. Also, the RDMs of the raw pixel data and the target labels are shown.

B MCNEMAR TEST P-VALUES

Table 1: P-values for the exact McNemar test for the pairwise comparison of the baseline, baseline with dropout (Dropout), baseline with soft-targets (Soft), baseline with RDL, and teacher networks on MNIST. An * indicates that models are significantly different using a threshold of 0.05.

Models	Baseline	Dropout	Soft	RDL	Teacher
Baseline	-	0.094	<0.001*	0.038*	<0.001*
Dropout	0.094	-	0.077	0.649	0.031*
Soft	<0.001*	0.077	-	0.161	0.880
RDL	0.038*	0.649	0.161	-	0.131
Teacher	<0.001*	0.031*	0.880	0.131	-

Table 2: P-values for the exact McNemar test for the pairwise comparison of the baseline, baseline with dropout (Dropout), baseline with soft-targets (Soft), baseline with RDL, and teacher networks on CIFAR-10. An * indicates that models are significantly different using a threshold of 0.05.

Models	Baseline	Dropout	Soft	RDL	Teacher
Baseline	-	<0.001*	0.106	<0.001*	<0.001*
Dropout	<0.001*	-	<0.001*	<0.001*	0.014*
Soft	0.106	<0.001*	-	0.072	<0.001*
RDL	<0.001*	<0.001*	0.072	-	<0.001*
Teacher	<0.001*	0.014*	<0.001*	<0.001*	-

Table 3: P-values for the exact McNemar test for the pairwise comparison of the baseline, baseline with fine-tuning, and baseline with RDL. An * indicates that models are significantly different using a threshold of 0.05.

Models	Baseline	Fine-Tuning	RDL
Baseline	-	0.192	<0.001*
Fine-Tuning	0.192	-	0.023*
RDL	<0.001*	0.023*	-